arXiv:1511.04150v1 [stat.ML] 13 Nov 2015

Deep Mean Maps

Junier B. Oliva* Dougal J. Sutherland* Barnabás Póczos Jeff Schneider
Carnegie Mellon University
{joliva,dsutherl,bapoczos,schneide}@cs.cmu.edu

Abstract

The use of distributions *and* high-level features from deep architecture *has become commonplace in modern computer vision. Both of these methodologies have separately achieved a great deal of success in many computer vision tasks. However, there has been little work attempting to leverage the power of these to methodologies jointly. To this end, this paper presents the Deep Mean Maps (DMMs) framework, a novel family of methods to non-parametrically represent distributions of features in convolutional neural network models.*

DMMs are able to both classify images using the distribution of top-level features, and to tune the top-level features for performing this task. We show how to implement DMMs using a special mean map layer composed of typical CNN operations, making both forward and backward propagation simple.

We illustrate the efficacy of DMMs at analyzing distributional patterns in image data in a synthetic data experiment. We also show that we extending existing deep architectures with DMMs improves the performance of existing CNNs on several challenging real-world datasets.

1. Introduction

Modern computer vision has seen extensive use of two separate machine learning methodologies: deep architectures and learning on distributions. However, research has only just begun to consider the joint use of these two methods. In this paper we present deep mean maps (DMMs), a novel framework to jointly use and learn distributions of high-level features from deep architectures. We show that DMMs are easy to implement into existing deep learning infrastructure, and are able to improve the results of existing architectures on various datasets.

A substantial amount of work has been devoted to constructing distribution-based features for vision tasks. These distribution-based approaches are often patch-level histogram features such as bag of words (BoW) representations [22], where the "words" are local features such as histogram of gradients (HoG) [2] or dense SIFT features [25]. More recent developments have extended distribution-based methods beyond histograms to nonparametric continuous domains [30, 27]. Such methods are adept at providing robust representations that and give a holistic, aggregate view of an image helpful in complex classification tasks.

Deep architectures, which compose many layers of computational neural units, have recently achieved a tremendous amount of success in computer vision tasks by replacing hand-designed features with learned convolutional filters. In fact, deep architectures, often convolutional neural networks (CNNs), have almost become the de-facto standard in some datasets with their recent success [33]. It is believed that the depth and hierarchical nature of deep CNNs lead to top-level convolutional features that exhibit semantically meaningful representations [20]. Typically, the last levels of these features are then concatenated and fed through a few fully-connected layers to obtain a final classification result.

It would be informative to study the distributions of these high-level features in an image for supervised tasks (e.g. scene classification), rather than treating the extracted features simply as a vector. In this paper we address the lack of the application of distributions in deep architectures through the development of a deep mean map layer, which will provide a fast and scalable featurization of top-level convolutional features that will be representative of their distributions in a nonparametric fashion.

This combination of distribution-based and deep architecture approaches has seen some initial success with an ad-hoc method of extracting features from distributions of fixed, pre-trained high-level features [39]. However, there has yet to been work that jointly learns high-level features and uses their distributions for supervised learning tasks. To this aim we present deep mean maps, which do so in a scalable, nonparametric fashion.

Deep mean maps employ mean map embeddings [7, 34, 35], which with the use of random features [31], provide a finite-dimensional nonparametric representation of distribu-

*These two authors contributed equally.

tions of top-level features for discriminant learning. We will show that deep mean maps can be implemented with typical CNN machinery, making both forward and backward propagation tractable for learning effective features and discriminant models based on their distributions.

Outline The rest of the paper is structured as follows. First, we detail our DMM framework in Section 2. We then study its behavior on a synthetic problem and illustrate its efficacy on several real-world datasets in Section 3. Section 4 then discusses the relationship to other method and gives concluding remarks.

2. Method

In this section we explicate our framework for learning high-level descriptive features, whose distributions are used in supervised image classification problems.

We focus on deep architectures of convolutional networks, which we briefly discuss below; however, our framework extends to other architectures. We show how to treat the distributions of high-level features of convolutional networks as inputs to aid in learning tasks via mean maps [34]. Furthermore, we show how to learn supervised tasks using mean maps in a manner that scales to large datasets via random Fourier features [31]. Lastly, we demonstrate that one can incorporate mean maps into a convolutional network to perform supervised tasks whilst tuning the network's high-level features for this purpose.

2.1. High-level features from deep architectures

The use of deep architectures for learning high-level descriptive features has been extensively explored over the last decade [9]. It is believed that the deep architectures present in many modern models are instrumental to learning features for both supervised [16, 14] and unsupervised [21, 15] tasks. While such high-level features may be extracted in a variety of different deep architectures, we focus on deep convolution neural networks (CNNs) [17, 18, 19, 16].

CNNs have been used with great success in many vision tasks. Inspired by the study of visual cortices [17], CNNs use successive applications of locally-receptive convolution, subsampling, nonlinear transformation, and normalization layers. Through their localized nature, CNNs enjoy a reduction in parameters over general neural networks and may exploit the spatial structures present in image data. Furthermore, with the use of modern non-linear units such as ReLUs [28], efficient GPU implementations, and large datasets, modern CNNs have been able to employ more layers to achieve state of the art performance [14, 37]. The high accuracy and spatial structure of deep CNN models make the outputs of top-level convolution layers excellent high-level features for image classification. Thus, we look to leverage the distribution of these features to perform supervised tasks.

2.2. Scalable mean map embeddings

The main tool that we will use to learn with distributions of high-level features are mean map embeddings, which we detail below. We show how mean map embeddings, along with random features [31], can be implemented using typical CNN layers for scalable learning of distributions of high-level features in vision tasks.

2.2.1 Mean map embeddings

Mean map embeddings (MMEs) of distributions serve as a way to embed distributions into a reproducing kernel Hilbert space (RKHS) [7, 34], allowing one to apply many familiar kernel methods, such as SVMs and GPs, to datasets of distributions through the kernel trick. Formally, the MME of a distribution P in the RKHS induced by a kernel $K(x, y) = \langle \phi(X), \phi(y) \rangle_{\mathcal{H}}$ is defined as:

$$\mu_P \equiv \mathbb{E}_{X \sim P} \left[K(X, \cdot) \right] = \mathbb{E}_{X \sim P} \left[\phi(X) \right]. \tag{1}$$

If K is a *characteristic* kernel (as is the RBF kernel), then μ_P uniquely identifies P [35]. Moreover, one may apply an element of the RKHS $f \in \mathcal{H}$ to μ_P by:

$$\langle f, \mu_P \rangle_{\mathcal{H}} = \mathbb{E}_{X \sim P} \left[\langle f, K(X, \cdot) \rangle_{\mathcal{H}} \right] = \mathbb{E}_{X \sim P} \left[f(X) \right]. \tag{2}$$

Thus, given that we may compute the inner-product between the embedding of a distribution P in the RKHS, μ_P, and another element of the RKHS, f, we may set up a loss w.r.t. (2) to learn over distributions.

2.2.2 Sample estimates

Of course, the calculation of μ_P necessitates the population mean of $\mathbb{E}_{X \sim P} \left[K(X, \cdot) \right]$, which will not be available in real-world data. Instead, one typically computes a finite sample estimate of μ_P given a sample $\mathcal{X} = \{X_j\}_{j=1}^n \overset{\text{iid}}{\sim} P$:

$$\mu_{\mathcal{X}} = \frac{1}{n} \sum_{j=1}^{n} K(X_j, \cdot). \tag{3}$$

Furthermore, one may compute $\langle f, \mu_{\mathcal{X}} \rangle_{\mathcal{H}}$ as:

$$\langle f, \mu_{\mathcal{X}} \rangle_{\mathcal{H}} = \frac{1}{n} \sum_{j=1}^{n} f(X_j). \tag{4}$$

Hence, we may apply kernel methods to distributions $\{P_i\}_{i=1}^N$ when given a data-set of sample sets $\mathcal{D} = \{\mathcal{X}_i\}_{i=1}^N$, with $\mathcal{X}_i = \{X_{ij}\}_{j=1}^{n_i}$.

2.2.3 Random feature embeddings

Unfortunately, due to the computation of a $N \times N$ Gram matrix of pairwise kernel evaluations, the application of kernel methods to the dataset $\mathcal{D} = \{\mathcal{X}_i\}_{i=1}^N$ will scale as $\Omega(N^2)$ — intractable for modern image datasets containing millions of images. To mitigate this, one may use the mean embedding of random features [31, 32] to estimate kernel evaluations in an approximate finite primal space. That is, for a shift-invariant kernel $K(x, y) = k(x - y)$ with $k : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ and $k(0) = 1$ we have that

$$K(x, y) \approx z(x)^T z(y) \quad \text{where} \tag{5}$$

$$z(x) = \tfrac{1}{\sqrt{2D}} \left(\cos(\omega_1^T x + b_1), \ldots, \cos(\omega_D^T x + b_D)\right)^T, \tag{6}$$

$b_d \overset{\text{iid}}{\sim} \text{Unif}[0, 2\pi]$, and $\omega_d \overset{\text{iid}}{\sim} \rho$, where ρ is the spectral distribution of k (i.e. its Fourier transform). Thus, we can approximate the kernel evaluation $K(x, y)$ using a finite dimensional dot-product, even for characteristic kernels induced by an infinite feature map like the RBF kernel. For the RBF kernel,

$$K(x, y) \equiv \exp\left(-\frac{1}{2\sigma^2} \|x - y\|^2\right),$$

we have that $\rho(\omega) = \mathcal{N}(\omega; 0, \sigma^{-2} I)$. Furthermore, using the random feature map $z(\cdot)$ we can approximate the mean map μ_P as:

$$\hat{\mu}_{\mathcal{X}_i} \equiv \frac{1}{n} \sum_{j=1}^{n} z(X_j) \approx \mathbb{E}_{X \sim P}\left[\phi(X)\right] = \mu_P. \tag{7}$$

Moreover, we can approximate the inner product of μ_{P_i} and $f \in \mathcal{H}$, $f(x) = \sum_{l=1}^{\infty} \alpha_l K(x_l, x)$ as:

$$\langle f, \mu_{P_i} \rangle_{\mathcal{H}} \approx \frac{1}{n_i} \sum_{j=1}^{n_i} f(X_j) \tag{8}$$

$$= \frac{1}{n_i} \sum_{j=1}^{n_i} \sum_{l=1}^{\infty} \alpha_l K(X_j, x_l) \tag{9}$$

$$= \frac{1}{n_i} \sum_{j=1}^{n_i} \sum_{l=1}^{\infty} \alpha_l z(X_j)^T z(x_l) \tag{10}$$

$$= \left(\frac{1}{n_i} \sum_{j=1}^{n_i} z(X_j)\right)^T \left(\sum_{l=1}^{\infty} \alpha_l z(x_l)\right) \tag{11}$$

$$= \hat{\mu}_{\mathcal{X}_i}^T \psi, \tag{12}$$

where $\psi = \sum_{l=1}^{\infty} \alpha_l z(x_l)$. Hence, we see that one can estimate RKHS inner products of mean maps as finite-dimensional dot products in a primal space induced by the random features $z(\cdot)$. Thus, we may avoid the costly $\Omega(N^2)$ scaling by working in this approximate primal space. Below we expand on how to learn high-level features whose distributions are used in supervised tasks through mean map embeddings of random features.

2.3. MMEs of CNN features

Given that CNNs have proven to be successful feature extractors for classification tasks, it is only natural that we look to study the distributions of said features across spatial dimensions. For many vision tasks, such as scene classification, the distribution of filter values will be highly informative, since we expect the presence of features across various locations to be relevant for such supervised tasks. In fact, [39] achieved state-of-the-art performance whilst only considering fixed, pre-trained simple parametric approximations of the distribution of high-level features through ad-hoc summary statistics. In a similar vein, [26] uses a final global average pooling layer on convolution features that "sums out the spatial information," "is more robust to spatial translations," and generalizes better. One can view this global average pooling as representing the distribution of features by its mean; our method allows richer (and, in fact, *unique*) representations of these distributions. We now illustrate how to include a layer that computes the mean map embedding of convolution filters in order to both learn informative features and use their distributions to discriminate in supervised tasks.

Suppose that the top convolution layer of a CNN (after any subsampling) produces a tensor $C_i \in \mathbb{R}^{m \times h \times w}$ for the i^{th} input image fed forward through the network, where here m is the number of convolutional filters, and h, w are spatial dimensions of super pixels (see Figure 1). We wish to consider the distributions of the $m \times 1$ vectors of filter values:

$$\mathcal{X}_i = \{C_{i,j,l} \in \mathbb{R}^m \,|\, 1 \leq j \leq h,\ 1 \leq l \leq w\}, \tag{13}$$

where $C_{i,j,l} = (C_i(1, j, l), \ldots, C_i(m, j, l))^T$. To do so, we will calculate the MME of $\mathcal{X}_i$ (1) in what we term as the mean map layer. The mean map layer calculates $\hat{\mu}_{\mathcal{X}_i}$ and propagates it forward in the network, allowing one to classify images with distributions of high-level features that are tuned for this distributional representation.

We now show that the mean map layer can be represented in terms of standard CNN mechanisms. Specifically, the mean map layer performs: D scaled 1×1 convolutions that are biased by uniformly random offsets, a cosine nonlinearity layer, and global-average pooling. One can see that $\hat{\mu}_{\mathcal{X}_i}$ (7) is computed as the mean of the set of vectors $\{z(C_{i,j,l}) \,|\, 1 \leq j \leq h,\ 1 \leq l \leq w\}$. Let $Z_i \in \mathbb{R}^{D \times h \times w}$ be the tensor where

$$Z_i(d, j, l) = [z(C_{i,j,l})]_d = \cos\left(\omega_d^T C_{i,j,l} + b_d\right), \tag{14}$$

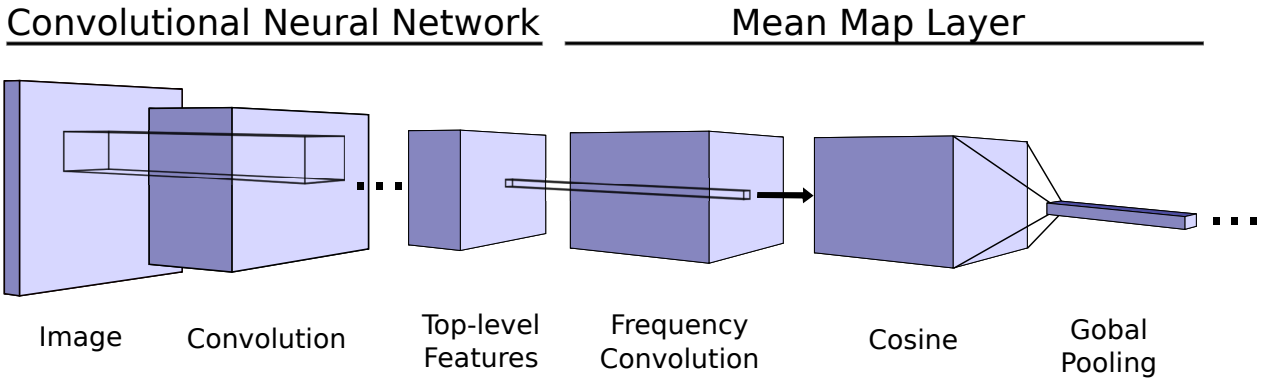

Figure 1. CNN with a mean map layer. We consider the top-level features resulting from several convolution/pooling layers. To compute the mean map embedding to the filter vectors , we perform the following: 1) take 1×1 convolution with frequencies $\{\omega_d\}_{d=1}^D$; 2) take element-wise cosines; 3) do global average pooling to compute a D-length vector.

and $\{b_d\}_{d=1}^D \overset{\text{iid}}{\sim} \text{Unif}[0, 2\pi]$ are drawn before training time and are held fixed throughout. Note that we have dropped the constant $\frac{1}{\sqrt{2D}}$ since it may be factored into linear operations above in the network. To compute these vectors, and their mean we perform (see Figure 1):

Frequency Convolution The term $\omega_d^T C_{i,j,l} + b_d$ in (14) is just a 1×1 convolution. Hence, we perform D 1×1 convolutions (one for each frequency ω_d).

Cosine Z_i is computed via a cosine layer that computes element-wise cosines to the convolutions above.

Global Pooling The computation of the d^{th} element of $\hat{\mu}_{\mathcal{X}_i}$ is just the average:

$$[\hat{\mu}_{\mathcal{X}_i}]_d = \frac{1}{hw} \sum_{j=1}^{h} \sum_{l=1}^{w} Z_i(d, j, l). \tag{15}$$

I.e. $\hat{\mu}_{\mathcal{X}_i}$ is computed via global average pooling of the channels in Z_i.

Hence, it is clear that the mean map layer is efficient to implement and train in terms of existing CNN operations and layers.

Note that in practice one may learn the kernel used in the mean map by learning frequencies $\{\omega_d\}_{d=1}^D$, or one may consider a fixed kernel by drawing $\{\omega_d\}_{d=1}^D \overset{\text{iid}}{\sim} \rho$ (where ρ is the spectral density for the kernel, e.g. $\mathcal{N}(0, \sigma^{-2} I)$ for the RBF kernel). Furthermore, one may learn the scale for a fixed kernel; e.g. one may learn the bandwidth of the RBF kernel in the MME by drawing $\{\omega_d\}_{d=1}^D \overset{\text{iid}}{\sim} \mathcal{N}(0, I)$ and parameterizing the cosine layer in terms of $a \in \mathbb{R}$ as:

$$Z_i(d, j, l) = \cos\left(e^a \omega_d^T C_{i,j,l} + b_d\right), \tag{16}$$

where here $\sigma^2 = e^{2a}$, and can be learned via back-propagation.

After being fed through the MME layer (with the global-average pooling of random features), the i^{th} image in a batch will have the mean map embedding of its top-level convolution features, $\hat{\mu}_{\mathcal{X}_i} \in \mathbb{R}^D$, computed. These vectors can then then be fed through the network in a fully connected fashion to perform learning w.r.t. a loss. Note that linear operations of the MME vectors, $\hat{\mu}_{\mathcal{X}_i}^T \psi$ where $\psi \in \mathbb{R}^D$, may be interpreted as an RKHS inner-product with an element $f \in \mathcal{H}$ (12). Moreover, it is straightforward to perform back-propagation through the MME layer by calculating derivatives of (14, 16) w.r.t. inputs and parameters.

3. Empirical Evaluation

We demonstrate the effectiveness of our approach on two image classification problems. All experiments are implemented in the Caffe deep learning framework [11]; code will be made publicly available.

3.1. Synthetic Data

Below we show the ability of the mean map layer to uncover visual distributional patterns with a simple synthetic data example.

We generate data using a bank of $M = 112$ texture images found in the Colored Brodatz Texture (CBT) database [1]. Each class represents a distribution of patches in the generated images over the texture bank. We consider a dataset of 8 such classes, generated at random. Specifically, the i^{th} class corresponds to a mixture of 30 Dirichlet distributions, P_i, where each mixture component has support in the M-dimensional simplex (one dimension for each texture image). Given the Dirichlet mixture for the class, we generate a 120×120 image as follows: pick a random location for a 12×12 patch uniformly at random, draw $\pi \sim P_i$, pick a random patch from each image $\{p_j\}_{j=1}^M$; set the patch to $p = \sum_{j=1}^M \pi_j p_j$, overlay on chosen location, repeat 1500

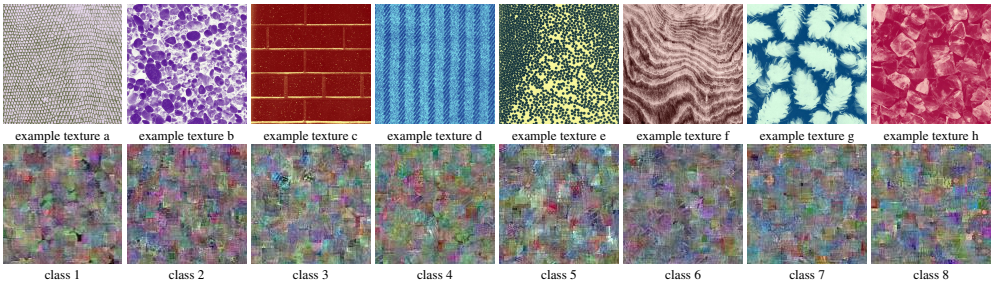

Figure 2. Top row: example images from the CBT database. Bottom row: random synthetic data images from the 8 classes.

times. We generated a set of 800 total training images (100 per class) and 4000 test images (500 per-class).

We used a simple CNN network with the following structure: a convolution of 58 filters with kernel sized 12×12 and a stride of 1; a ReLU layer; 12×12 max pooling with a stride of 6. We appended a mean map layer to the end of the network (`mml`). We also compared to appending two hidden layers each of 4096 nodes with ReLU non-linearity (`hid`), or directly classifying on the max pooled features (`lin`). Each network was trained with SGD.

We look at the test accuracies of the aforementioned three networks on a subset of 80 training images (10 from each class) and all 800 (see Figure 3). It is interesting to note that even with only 10 instances from each class `mml` is able to almost perfectly distinguish the classes, and outperforms the other networks by a large margin. This is because the `mml` has the built in mechanisms to learn over distributions of images features and thus may do so in a sample efficient fashion.

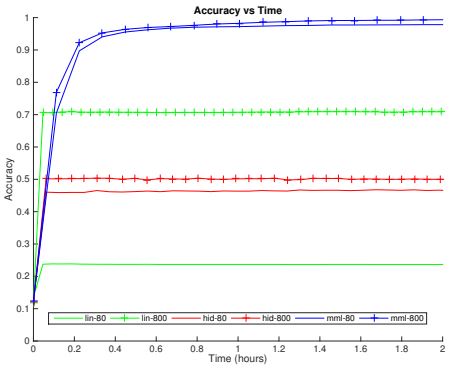

Figure 3. Accuracy vs. time for synthetic data experiment. Blue curves correspond to the accuracies for the `mml` network with the mean map layer. Green curves correspond to the `lin` network, which feeds max-pooled filters directly to a classifier. Lastly, red curves correspond to the `hid` network, which uses two hidden layers to classify. Solid and crossed lines indicate a training set of 80 and 800 respectively.

3.2. Adaptation

In this section we illustrate the effectiveness of deep mean maps on several real-world scene classification datasets. We shall begin with a network trained on the ImageNet dataset [33] and fine-tune it for each dataset.

3.2.1 Datasets

Flickr Style The first dataset we consider, known as Flickr Style, was assembled by [13]. It contains 80,000 images in 20 style categories defined by user groups such as "Bright," "Noir," "Long Exposure", "Geometric," and "Romantic." We use the standard 80-20 train-test split, and further divide the training images with an 80-20 train-validation split for model selection. Figure 4 shows an image from each class.

Wikipaintings This dataset, also assembled by [13], contains about 83,000 images in 25 genres including "Impressionism," "Baroque," "Surrealism," and "Cubism." Each category contains at least 1,000 images. We divide the dataset randomly into 53,140 training, 13,285 validation, and 16,612 test images. Figure 5 shows an image from each class.

Places The Places-205 dataset [40] contains 2,448,873 training, 20,500 validation, and 41,000 test images in 205 scene categories, including "runway," "ski resort," "trench," "kasbah," and "dorm room." The test labels are not public; submissions must be scored on a server run by the authors of [40]. Figure 6 shows sample images.

3.2.2 Networks

We consider both a minor variation of the model of [14] known as CaffeNet (`bvlc_reference_caffenet`), and the GoogLeNet model of [38] (`bvlc_googlenet`). We fine-tune each network with SGD, with a consistent learning rate schedule for each architecture, and pick the model of our occasional snapshots with highest validation accuracy.

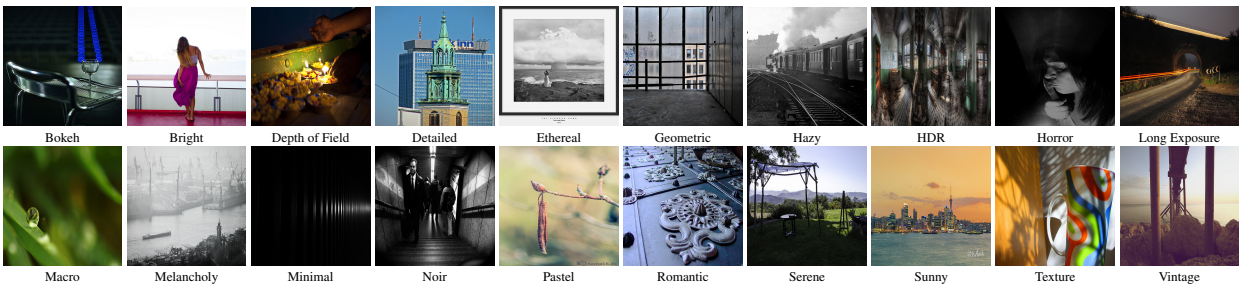

Figure 4. A random image from each of the 20 classes in Flickr Style.

Figure 5. One random image from each of the 25 classes in Wikipaintings.

Training time for each variant is comparable to the original network.

We extend these networks to use mean map layers (with 4,096 frequencies) in the following ways:

Replacing By replacing the fully connected layers with a mean map layer after the top-level features of the CNN, we look to classify solely on the mean map embedding of these top level features.

Replicating One may send the top-level features both to a hidden layer as well as to a mean map layer; this allows one to classify with a combination of the vectorized top-level features and the mean map embedding.

Forking Furthermore, rather than sending the same top-level features into the hidden and mean map layers above as in replication, one may learn two sets of distinct top-level-features to be fed to a hidden and mean map layer respectively.

We also consider the following variants on each model:

Dropout Add a dropout layer after the mean map embedding. This may encourage the network to avoid relying on any one frequency in the random embedding, thereby more faithfully replicating the kernel.

Hidden Add a fully-connected hidden layer between the mean map embedding and the classifier.

Frequency Learning Optimize the individual frequencies sampled, rather than just their scale.

Model selection was performed on the validation set; the three best MME models, along with the original model, were each scored on the test set.

3.2.3 Results

Tables 1-3 present adaptation results for each of the three datasets for both CaffeNet and GoogleNet with the three best DMM extensions (DMM type) on the validation set and the "base" model (DMM type "None").

The best model variant picked according to validation score nearly always outperform the equivalent "base" model on the test set on Flickr Style and Wikipaintings. Hence we see that DMMs, through the use on distributions of top-level features, are able top improve existing deep-architectures. Furthermore, while the flexible nature of the mean map layer allows it to extend existing networks in many different variations, DMMs seem to generalize well empirically. Thus, a simple validation set approach is effective for choosing how to extend a network's architecture with DMMs.

4. Conclusion

4.1. Related Work

As previously mentioned, deep architectures for learning high-level features has been explored in a myriad of work [9, 16, 14, 21, 15]. While such high-level features may be extracted in a variety of different deep architectures, we focus on deep convolution neural networks (CNNs) [17, 18, 19, 16]. Modern CNNs have been able to employ more layers [14, 37] without unsupervised pretraining thanks to their reduced parameterization (as compared to general neural networks), recent advents in training methods and nonlinear units [10, 28], and large datasets [33].

Also, a great deal work has been devoted to constructing distribution-based features for vision tasks. These

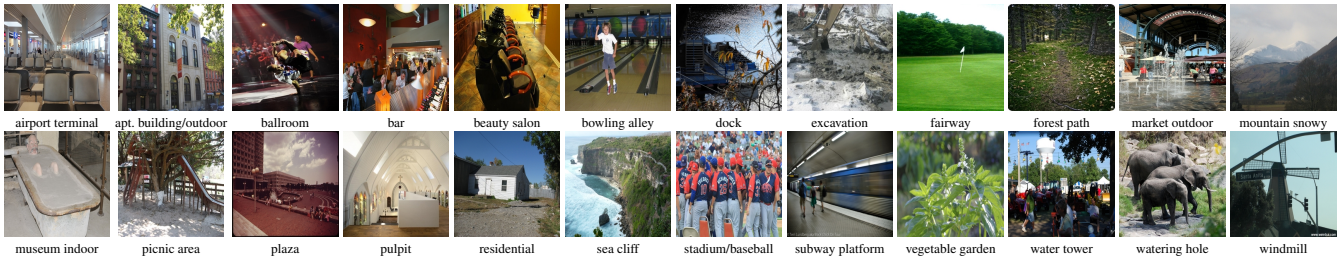

Figure 6. Random images from 24 of the 205 classes in Places-205.

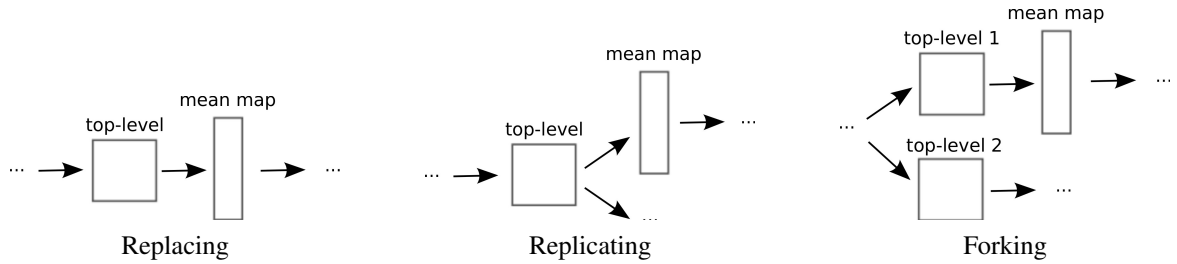

Figure 7. Illustration of each method for extending existing deep models. Top-level refers to the feature from the last convolutions layer. *Replacing* here we append a mean map layer to the top-level convolutional features. *Replicating* here we copy the top-level convolutional features and feed one copy to the mean map layer while passing the other copy through to classification layers. *Forking* here we make two distinct top-level convolutional features, one that is fed to a mean map layer and the other is fed through to classification layers.

distribution-based approaches are often histograms such as BoW representations [4, 22], HoG features [2], and dense-SIFT [25]. Moreover, there have been efforts to extended distribution-based methods to nonparametric continuous domains [30].

There has been work that explores the use of distributions of high-level deep architecture features with ad-hoc features from distributions of fixed, pre-trained high-level features [39]. In contrast, this paper learns high-level features and uses their distributions in a scalable, nonparametric fashion with mean maps.

Mean map embeddings [34, 7, 35] have served as a practical and flexible method to learn using datasets of sample sets drawn from distributions [27, 36]. Furthermore, the use of random features [31, 32] has allowed for scalable learning over distributions [29, 12, 24, 5]. Mean map embeddings have also been used extensively for the comparison and two-sample testing of distributions [7, 8, 6]. Recent research has also explored the use of mean maps in deep architectures for building fast sampling mechanisms, by conducting two-sample tests between generated sample points and original data [23, 3]; this work differs from ours in using mean maps in the loss function, rather than in the architecture itself.

4.2. Remarks

This paper presents DMMs, a novel framework to jointly use and learn distributions of top-level features in deep architectures. We showed that DMMs, though a mean map layer, may use the mean embedding of top-level features to non-parametrically represent the distributions of the features.

The use of mean embeddings with random features allow the mean map layer to non-parametrically represent distributions of top-level features whilst still scaling to large image datasets. Also, inner-products on the mean embeddings through the network are interpretable as RKHS inner products on the distributional embeddings, allowing one to build a strong theoretical foundation [36]. Furthermore, we showed that the mean map layer may be implemented using typical CNN operations, making both forward and backward propagation simple to do with DMMs.

Moreover, we illustrated the aptitude of the mean map layer at learning distributions of visual features for discrimination in a synthetic data experiment. We saw that even with very few instances the mean map layer allowed a network to quickly learn to distinguish visual distributional patterns in a sample efficient manner. Lastly, we showed that DMMs may be used to extend several existing state-of-the-art deep-architectures and improve their performance on various challenging real-world datasets. Indeed, the mean map layer prove flexible and capable of extending networks an a variety of ways; due to a propensity to generalize well, it is simple to choose an extension method with a straightforward validation set approach. Thus, it is clear that the DMM framework can successfully build on the myriad of state-of-the-art deep architectures available.

Architecture	DMM Type	Variants	Validation Top-1	Test Top-1	Test Top-3
CaffeNet	None	-	37.69%	36.98%	62.17%
CaffeNet	Replacing	Dropout, Hidden, Frequency Learning	**38.87%**	**38.61%**	**64.47%**
CaffeNet	Replacing	Dropout, Hidden	38.27%	37.91%	63.64%
CaffeNet	Replicating	Dropout, Hidden	38.27%	37.71%	63.31%
GoogLeNet	None	-	40.44%	39.93%	66.32%
GoogLeNet	Replicating	Dropout, Frequency Learning	**40.74%**	**40.20%**	**66.66%**
GoogLeNet	Replicating	Dropout	40.66%	40.10%	66.26%
GoogLeNet	Replicating	Dropout, Hidden, Frequency Learning	40.48%	39.66%	66.32%

Table 1. Adaptation results on the Flickr Style dataset. **Boldface** indicates best score for CaffeNet and GoogLeNet.

Architecture	DMM Type	Variants	Validation Top-1	Test Top-1	Test Top-3
CaffeNet	None	-	55.96%	54.94%	79.99%
CaffeNet	Replacing	-	**56.03%**	**55.16%**	**80.63%**
CaffeNet	Forking	Dropout	55.92%	55.03%	80.29%
CaffeNet	Forking	-	55.76%	54.91%	80.45%
GoogLeNet	None	-	**58.94%**	57.49%	83.72%
GoogLeNet	Replicating	Dropout	58.18%	57.06%	83.64%
GoogLeNet	Replicating	Dropout, Frequency Learning	57.96%	**57.56%**	**83.84%**
GoogLeNet	Replicating	Dropout, Hidden, Frequency Learning	57.70%	56.36%	83.25%

Table 2. Adaptation results on the Wikipaintings dataset. **Boldface** indicates best score for CaffeNet and GoogLeNet.

Architecture	DMM Type	Variants	Validation Top-1	Test Top-1	Test Top-5
CaffeNet	None	-	45.92%	45.69%	77.45%
CaffeNet	Forking	-	**46.86%**	**46.62%**	**77.99%**
CaffeNet	Replicating	-	46.40%	46.31%	77.97%
CaffeNet	Forking	Dropout	46.19%	46.27%	77.95%
GoogLeNet	None	-	47.50%	47.50%	79.83%
GoogLeNet	Forking	Dropout, Hidden	**49.17%**	48.99%	**80.45%**
GoogLeNet	Forking	Dropout, Hidden, Frequency Learning	48.80%	**49.00%**	80.42%
GoogLeNet	Forking	Dropout, Frequency Learning	48.32%	48.56%	80.21%

Table 3. Adaptation results on the Places dataset. **Boldface** indicates best score for CaffeNet and GoogLeNet.

References

[1] Colored brodatz texture database. http://multibandtexture.recherche.usherbrooke.ca/colored%20_brodatz.html. 4

[2] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893. IEEE, 2005. 1, 7

[3] G. K. Dziugaite, D. M. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. *arXiv:1505.03906*, 2015. 7

[4] L. Fei-Fei and P. Perona. A Bayesian hierarchical model for learning natural scene categories. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 524–531. IEEE, 2005. 7

[5] S. R. Flaxman, Y.-X. Wang, and A. J. Smola. Who supported Obama in 2012?: Ecological inference through distribution regression. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 289–298. ACM, 2015. 7

[6] K. Fukumizu, A. Gretton, X. Sun, and B. Schölkopf. Kernel measures of conditional dependence. In *NIPS*, volume 20, pages 489–496, 2007. 7

[7] A. Gretton, K. M. Borgwardt, M. Rasch, B. Schölkopf, and A. J. Smola. A kernel method for the two-sample-problem. In *Advances in neural information processing systems*, pages 513–520, 2006. 1, 2, 7

[8] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. 7

[9] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006. 2, 6

[10] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012. 6

[11] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv:1408.5093*, 2014. 4

[12] W. Jitkrittum, A. Gretton, N. Heess, S. Eslami, B. Lakshminarayanan, D. Sejdinovic, and Z. Szabó. Kernel-based just-in-time learning for passing expectation propagation messages. *arXiv:1503.02551*, 2015. 7

[13] S. Karayev, M. Trentacoste, H. Han, A. Agarwala, T. Darrell, A. Hertzmann, and H. Winnemoeller. Recognizing image style. *arXiv:1311.3715*, 2013. 5

[14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, pages 1097–1105, 2012. 2, 5, 6

[15] Q. V. Le. Building high-level features using large scale unsupervised learning. 2, 6

[16] Y. LeCun. Learning invariant feature hierarchies. In *Computer vision–ECCV 2012. Workshops and demonstrations*, pages 496–505. Springer, 2012. 2, 6

[17] Y. LeCun and Y. Bengio. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10), 1995. 2, 6

[18] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2, 6

[19] Y. LeCun, F. J. Huang, and L. Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 2, pages II–97. IEEE, 2004. 2, 6

[20] Y. LeCun and M. Ranzato. Deep learning tutorial. International Conference on Machine Learning, 2013. 1

[21] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 609–616. ACM, 2009. 2, 6

[22] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *International journal of computer vision*, 43(1):29–44, 2001. 1, 7

[23] Y. Li, K. Swersky, and R. Zemel. Generative moment matching networks. *arXiv:1502.02761*, 2015. 7

[24] D. Lopez-Paz, K. Muandet, B. Schölkopf, and I. Tolstikhin. Towards a learning theory of causation. *arXiv:1502.02398*, 2015. 7

[25] D. G. Lowe. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pages 1150–1157. Ieee, 1999. 1, 7

[26] S. Y. M. Lin, Q. Chen. Network in network. 3

[27] K. Muandet, K. Fukumizu, F. Dinuzzo, and B. Schölkopf. Learning from distributions via support measure machines. *Advances in neural information processing systems*, pages 10–18, 2012. 1, 7

[28] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 807–814, 2010. 2, 6

[29] J. B. Oliva, W. Neiswanger, B. Póczos, J. Schneider, and E. Xing. Fast distribution to real regression. *arXiv preprint arXiv:1311.2236*, 2013. 7

[30] B. Póczos, L. Xiong, D. J. Sutherland, and J. Schneider. Nonparametric kernel estimators for image classification. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2989–2996. IEEE, 2012. 1, 7

[31] A. Rahimi and B. Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, pages 1177–1184, 2007. 1, 2, 3, 7

[32] A. Rahimi and B. Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Advances in neural information processing systems*, pages 1313–1320, 2009. 3, 7

[33] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, pages 1–42, 2014. 1, 5, 6

[34] A. Smola, A. Gretton, L. Song, and B. Schölkopf. A hilbert space embedding for distributions. In *Algorithmic Learning Theory*, pages 13–31. Springer, 2007. 1, 2, 7

[35] L. Song, K. Fukumizu, E. Gordon, A. Gretton, S. Hong, B. Schölkopf, E. Smola, V. Vishwanasan, and L. Williams. Learning via hilbert space embedding of distributions. 2008. 1, 2, 7

[36] Z. Szabó, A. Gretton, B. Póczos, and B. Sriperumbudur. Learning theory for distribution regression. *arXiv preprint arXiv:1411.2066*, 2014. 7

[37] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *arXiv preprint arXiv:1409.4842*. 2, 6

[38] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *arXiv:1409.4842*, 2014. 5

[39] J. Wu, B.-B. Gao, and G. Liu. Visual recognition using directional distribution distance. *arXiv preprint arXiv:1504.04792*, 2015. 1, 3, 7

[40] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. In *Advances in Neural Information Processing Systems*, 2014. 5